

PCT

WORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau



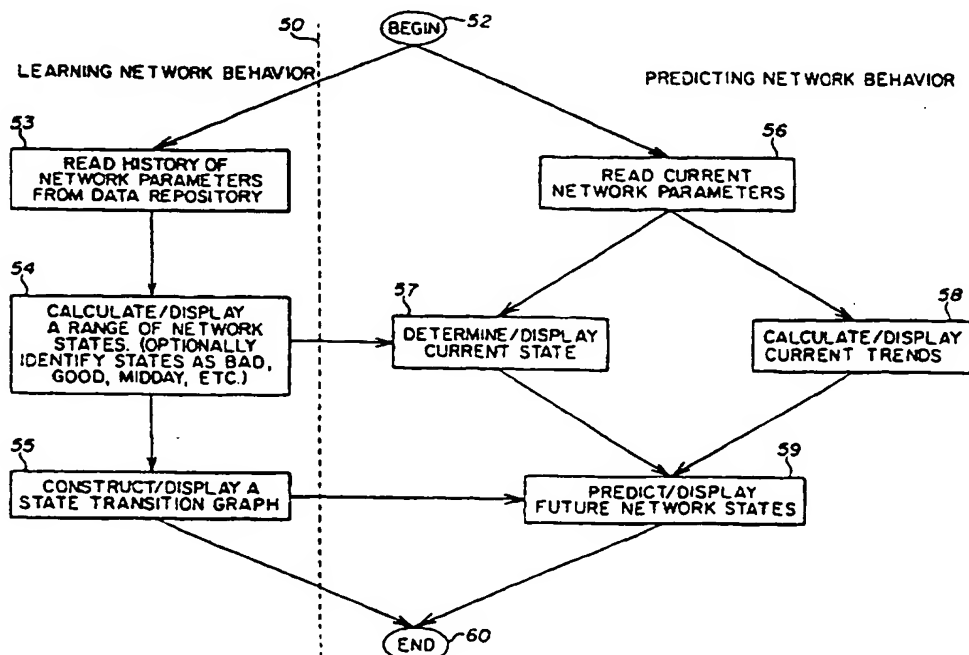
INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification 6 : H04L 12/24		A2	(11) International Publication Number: WO 96/24210
			(43) International Publication Date: 8 August 1996 (08.08.96)
(21) International Application Number: PCT/US96/01137		(81) Designated States: AL, AM, AT, AU, AZ, BB, BG, BR, BY, CA, CH, CN, CZ, DE, DK, EE, ES, FI, GB, GE, HU, IS, JP, KE, KG, KP, KR, KZ, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, TJ, TM, TR, TT, UA, UG, UZ, VN, ARIPO patent (KE, LS, MW, SD, SZ, UG), Eurasian patent (AZ, BY, KG, KZ, RU, TJ, TM), European patent (AT, BE, CH, DE, DK, ES, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, ML, MR, NE, SN, TD, TG).	
(22) International Filing Date: 24 January 1996 (24.01.96)			
(30) Priority Data: 08/382,666 2 February 1995 (02.02.95) US			
(71) Applicant: CABLETRON SYSTEMS, INC. [US/US]; 35 Industrial Way, Rochester, NH 03867 (US).			
(72) Inventors: LEWIS, Lundy; 480 Greenville Road, Mason, NH 03048 (US). DATTA, Utpal; 52 Pinecrest Drive, Bedford, NH 03110 (US).			
(74) Agent: HENDRICKS, Therese, A.; Wolf, Greenfield & Sacks, P.C., 600 Atlantic Avenue, Boston, MA 02210 (US).		Published Without international search report and to be republished upon receipt of that report.	

(54) Title: METHOD AND APPARATUS FOR LEARNING NETWORK BEHAVIOR TRENDS AND PREDICTING FUTURE BEHAVIOR OF COMMUNICATIONS NETWORKS

(57) Abstract

Apparatus and method for learning current network behavior and predicting future behavior which utilizes a state transition graph (40). The graph includes nodes (41-46) which represent network states, and arcs (48) which represent trends in observable network parameters that result in a transition from a current state to another state. For example, a watch service may be instituted on multiple ports of a router (3, 4), and the observed network traffic on the ports over time may be transformed into a state transition graph that represents network behavior. The network states may be labeled such as "good", or "bad", etc., according to a predetermined performance criteria. Once a state transition graph is constructed, the system may then monitor the current state and current trends of the network parameters in order to predict and display future network states. The system may include an automatic warning signal for alerting a user that the network is headed in the direction of a problematic state.



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AM	Armenia	GB	United Kingdom	MW	Malawi
AT	Austria	GE	Georgia	MX	Mexico
AU	Australia	GN	Guinea	NE	Niger
BB	Barbados	GR	Greece	NL	Netherlands
BE	Belgium	HU	Hungary	NO	Norway
BF	Burkina Faso	IE	Ireland	NZ	New Zealand
BG	Bulgaria	IT	Italy	PL	Poland
BJ	Benin	JP	Japan	PT	Portugal
BR	Brazil	KE	Kenya	RO	Romania
BY	Belarus	KG	Kyrgyzstan	RU	Russian Federation
CA	Canada	KP	Democratic People's Republic of Korea	SD	Sudan
CF	Central African Republic	KR	Republic of Korea	SE	Sweden
CG	Congo	KZ	Kazakhstan	SG	Singapore
CH	Switzerland	LJ	Liechtenstein	SI	Slovenia
CI	Côte d'Ivoire	LK	Sri Lanka	SK	Slovakia
CM	Cameroon	LR	Liberia	SN	Senegal
CN	China	LT	Lithuania	SZ	Swaziland
CS	Czechoslovakia	LU	Luxembourg	TD	Chad
CZ	Czech Republic	LV	Latvia	TG	Togo
DE	Germany	MC	Monaco	TJ	Tajikistan
DK	Denmark	MD	Republic of Moldova	TT	Trinidad and Tobago
EE	Estonia	MG	Madagascar	UA	Ukraine
ES	Spain	ML	Mali	UG	Uganda
FI	Finland	MN	Mongolia	US	United States of America
FR	France	MR	Mauritania	UZ	Uzbekistan
GA	Gabon			VN	Viet Nam

**METHOD AND APPARATUS FOR LEARNING NETWORK BEHAVIOR
TRENDS AND PREDICTING FUTURE BEHAVIOR
OF COMMUNICATIONS NETWORKS**

5

Field of the Invention

The present invention is directed to the monitoring, analysis and prediction of network behavior, and more specifically to an apparatus and method for learning and displaying a current behavioral state and possible transitions to other states and for monitoring current
10 behavior trends in order to predict future states of the network.

Background of the Invention

In the prior art, an averaging/thresholding method has been applied for learning the behavior trends of a network. Basically, this method observes the traffic in a network
15 segment over a period of time, for example a month, in order to determine an average or norm of the behavior over time. For example, the average bandwidth utilization in a network segment backbone may be represented by an interval $[x, y]$, where x is the lower threshold of utilization and y is the upper threshold. During a period of one month, the bandwidth utilization may for instance fall within an interval $[25, 40]$, with only a few stray values falling outside the interval,
20 where 25 is the lower threshold, and 40 is the upper threshold.

Commercial tools that implement the averaging method generally record bandwidth utilization data for some period of time and then use a statistical algorithm to calculate the norm. A current value of bandwidth utilization is then compared with the calculated norm. If the current value is outside the norm, an alarm is issued to warn the network
25 administrator of the discrepancy.

The averaging/thresholding method can be extended to find norms for traffic occurring within multiple network segments, in trunks that connect segments, and between individual nodes in a network. An example of a commercial tool that uses the averaging/thresholding method is the HP Network Advisor, sold by Hewlett Packard Company.
30 4 Choke Cherry Road, Rockville, MD 20850.

The averaging/thresholding method is useful for setting watches on network segments and for alerting the network administrator when the current traffic on a particular network segment exceeds a threshold. However, if the network administrator needs to understand the overall behavior patterns of the network, the averaging/thresholding method does
35 not suffice. A method that predicts network behavior would be extremely helpful in providing

the administrator more opportunity and time to intervene when the network appears to be moving toward a problematic state. With such a method, the administrator would be able to intervene before the problematic state occurred.

5

Summary of the Invention

The present invention is directed to a method and apparatus for learning behavior patterns of a network, for producing a state transition graph to represent the behavior patterns, and for predicting future states that the network may enter from a current state. This allows a network administrator to understand the behavior patterns of the overall network. In addition, by
10 being able to predict future states, it provides the network administrator with more opportunity and time to intervene whenever the network appears to be moving toward a problematic state.

The method of this invention determines a range of possible network states and transitions and displays the same in a state transition graph. The possible states and transitions are derived from a history of network parameters accumulated in the data repository of a network
15 management system. The current network parameters are then read for calculating the current state, current trends and the possible next states. Still further, by characterizing the next network states according to a performance criterion, the present invention can identify problematic states and warn the network administrator. Finally, the present invention displays the derived information about the network behavior to a user.

20 These and other advantages of the present invention are more particularly described in the following detailed description and drawings.

Brief Description of the Drawings

Fig. 1 is a schematic illustration of an apparatus for learning, predicting, and
25 displaying network behavior patterns according to this invention.

Fig. 2 is a table of representative network parameter data accumulated in a data repository of a network management system.

Fig. 3 shows various types of prior art displays for graphically illustrating network parameters.

30 Fig. 4 is a high-level topology view of a sample network containing multiple subnets and routers.

Fig. 5 is an example of a state transition graph generated according to the present invention.

Fig. 6 is a further exemplary display of a state transition graph with the current state and possible transitions from the current state all highlighted in bold.

5 Fig. 7 is a simplified state transition graph with only two states.

Fig. 8 is a flow diagram illustrating the method steps of the present invention.

Detailed Description

10 Fig. 1 is a block diagram illustrating the method and apparatus of the present invention. A network management system 14 monitors a live network 10 via communication link 12 over a period of time and passes the resulting network parameter data to a data repository 18 via communication link 16. Over a period of time, streams of data (e.g., of the type shown in Fig. 2) accumulate in the data repository 18. Data processor 22 accesses this accumulated data from the data repository 18 via communication link 20.

15 Within the data processor 22, method₁, method₂, ..., method_n (collectively labeled 24) refer to existing methods of transforming the data in data repository 18 into view₁, view₂, ..., view_n (collectively labeled 34) on graphical interface 32, via communication links 28. Fig. 3 shows examples of such prior art views, which may be pie graphs, bar graphs, and two dimensional graphs derived from the network parameters.

20 The present invention comprises both a new method and a new view, which are shown in Fig. 1 as a "Method to Determine Network Behavior Patterns" 26 (within processor 22) and a "View of Network Behavior Patterns" 36 (within graphical display 32 via communication link 30). The new method solves the following problem: Given a current state of a network and an observation of key network parameters, what will be the next state? In order to solve this
25 problem, we now define: the concept of a network state; examples of observable network parameters; and rules for predicting a transition from one state to another given the values for the network parameters.

30 Fig. 4 shows an exemplary high-level topology view of a network containing multiple subnets and routers (as an example of live network 10 in Fig. 1). The group of twelve icons 1 on the left of the figure, and the group of sixteen icons 2 on the right, may be Ethernet subnets. A pair of routers 3, 4 near the center are connected to these two groups of subnets, respectively, and the routers themselves are connected to each other via a coupler 5. Note that

other subnets and routers (not shown in Fig. 4) have trunks attached to the network via couplers 6 and 7, which in turn are connected to router 4.

For purposes of illustration, consider the leftmost router 3 in Fig. 4 and its twelve subnet group. Router 3 has twelve ports with traffic flowing to and from the respective twelve subnets. In order to monitor the amount of traffic on the network, watches may be placed on each of the twelve ports. These watches can measure the percentage of port utilization during a predetermined time interval.

The measurements may be recorded in a table, such as that shown in Fig. 2, and stored in the data repository 18. Fig. 2 shows a measured level of traffic on each of twelve ports P1, P2, ..., and P12, for a few minutes of time. Examples of network management systems (see 14 in Fig. 1) capable of implementing such a watch are: 1) Sniffer, Network General Corporation, 4200 Bohannon Drive, Menlo Park, CA 94025; 2) NetMetrix, Hewlett-Packard Corporation, 1 Tara Boulevard, Nashua, NH 03062; 3) LANalyzer, Novell, Inc., 122 East 1700 South, Provo, UT 84606-6194; and 4) Spectrum™ Network Management System, Cabletron Systems, Inc., Rochester, New Hampshire. In addition, the network management system 14 may include network management platforms, network monitors, or basic low-level programs such as "Etherfind" on Sun workstations, or "Netsnoop" on Silicon Graphics IRIX workstations.

Port utilization data measured for one month in the form of Fig. 2 would take up thousands of pages, and would not provide a general description of network behavior. The prior art methods allow transformation of the accumulated data into graphical views such as an x-y plot, pie graph, and bargraph (see Fig. 3.) In addition, the prior art averaging algorithm may be applied to the measured data in order to set alarm thresholds for each port. However, none of these methods allow a network administrator to understand overall patterns in network behavior or to predict network behavior.

The method of the present invention transforms the numeric data in Fig. 2 into a state transition diagram (a.k.a. a deterministic finite automaton) that represents the overall behavior patterns of the network. Fig. 5 shows an example of such a state transition graph 40, wherein nodes 41-46 represent six possible network states and arcs 48 connecting the nodes represent trends in observable network parameters that result in a transition from one state to another state. The trends are labeled with one or more port numbers (from Fig. 2) each followed by a symbol which indicates how the observed parameter is changing. The states are represented as nodes and embody the twelve parameters (where each parameter represents traffic at one of

the twelve ports). The value of each parameter may be an interval $[x, y]$, where x represents a lower bound of port utilization and y represents an upper bound, over a predetermined time period. Accordingly, a network state for the data of Fig. 2 may be represented as follows:

State S

$\{P1 = [x1, y1],$
 $P2 = [x2, y2],$
 $P3 = [x3, y3],$
 $P4 = [x4, y4],$
 $P5 = [x5, y5],$
 $P6 = [x6, y6],$
 $P7 = [x7, y7],$
 $P8 = [x8, y8],$
 $P9 = [x9, y9],$
 $P10 = [x10, y10],$
 $P11 = [x11, y11],$
 $P12 = [x12, y12] \}.$

The different states and trends are determined based on historic data and any one of the supervised or unsupervised learning methods described hereinafter.

The current numeric values on the ports P1, P2, ..., P12 are measured in short time increments, for example 10 minutes. A simple process may be used to translate the numeric values into symbolic values. For the state transition graph of Fig. 5. the language of the symbolic values that represent trends in current network parameters is as follows:

0 no traffic;
1 maximum traffic;
- moderately decreasing traffic;
-- quickly decreasing traffic;
= stable traffic;
+ moderately increasing traffic; and
++ quickly increasing traffic.

A prediction of future network behavior is determined from the state transition graph, the current state of the network, and the current network trends. For example, in the state transition graph of Fig. 5, if the network were in state 1 and a measurement of the current network parameters indicates that utilization of port 3 is increasing quickly, the present method would predict that the network will enter state 4. Upon transition to state 4, the network can either make a transition to state 3 or state 5, or return to state 1.

The present method labels the states according to a performance criterion. In Fig. 5, if state 3 were designated a "bad" state, then state 4 would be a potentially problematic state because state 4 could enter the bad state if the utilization on port 6 were to increase quickly. The present method warns the network administrator of potentially problematic states.

Finally, the present method may display all of the derived information concerning the network behavior patterns to the user, i.e., the state transition graph, the current state, and the possible future states. For example, Fig. 6 shows an exemplary display of a state transition graph where state 4 is highlighted as the current state. Also, the possible transitions via the arcs from state 4 are highlighted to emphasize the possible next states. In the alternative, the current state and the possible future states may be color coded on the display. In addition, labels corresponding to a predetermined performance criterion may be provided to the user.

The method of this invention has taken a history of network behavior as embodied for example in the data of Fig. 2 and has transformed it into a state transition diagram as in Fig. 5. This transformation is referred to generally as "learning." Several methods of learning exist, but the two main classes of learning being known as "supervised" and "unsupervised."

In supervised learning, labels are applied to the numeric data at each time increment. Labels may depict a performance criterion, and examples of such labels are "good," "bad," "state 1," and "evening state." A supervised learning algorithm generalizes over particular network parameters in order to produce a concept of a "good" state, or a "bad" state, according to the performance criterion.

In unsupervised learning, no labels are applied to the numeric data. This type of learning algorithm discovers and delineates key network states and

labels the states with arbitrary names such as "state 1" and "state 2." However, the network administrator may wish to re-label the states with more meaningful names.

Examples of prior art supervised and unsupervised learning methods include the following:

A. Supervised Methods

- (i) Iterative Dichotomizing Third (ID3) Algorithms;
- (ii) Multilayer Perceptrons (a.k.a. neural networks with backpropagation learning); and
- (iii) Recurrent Neural Networks.

B. Unsupervised Methods

- (i) Adaptive Resonance Theory (ART) Networks;
- (ii) Kohonen's Self Organizing Feature Maps; and
- (iii) Clustering Algorithms.

These methods differ with respect to efficiency, correctness, and

ease-of-implementation. The present invention is not tied to any particular one of these methods.

The following pseudo-code is for a particular multilayer perceptron learning embodiment that may be implemented in a C++ software program:

Definitions

$u_{l,j}$	output of the j th node in layer l
$w_{l,j,i}$	weight which connects the i th node in layer $l-1$ to the j th node in layer l
x_p	p th training sample
$u_{0,i}$	i th component of the input vector
$d_j(x_p)$	desired response of the j th output node for the p th training sample
N_l	number of nodes in layer l
L	number of layers
P	number of training patterns

Backpropagation Learning Algorithm**procedure BACK_PROP**

Initialize the weights to small random values;

repeat5 Choose next training pair (x,d) and let the 0th layer be $u_0=x$;

FEED_FORWARD;

COMPUTE_GRADIENT;

UPDATE_WEIGHTS;

until termination condition reached;10 **end;** {BACK_PROP}**subroutine FEED_FORWARD****for** layer=1 to L **do****for** node = 1 to N_{layer} **do**15 $u_{\text{layer,node}} = f(\sum_{i=0}^{N_{\text{layer-1}}} w_{\text{layer,node},i} u_{\text{layer-1},i});$ **endloop****endloop**20 **end;** {FEED_FORWARD}**subroutine COMPUTE_GRADIENT****for** layer = L to 1 **do****for** node = 1 to N_{layer} **do**25 **if** layer = L **then** $e_{L,node} = u_{L,node} - d_{node};$ **else** $e_{\text{layer,node}} =$ $\sum_{m=1}^{N_{\text{layer-1}}}$ $e_{\text{layer+1},m} u_{\text{layer+1},m} (1-u_{\text{layer+1},m}) w_{\text{layer+1},m,node};$

30

endloop**for all** weights in layer layer **do** $g_{\text{layer},j,i} = e_{\text{layer},j} u_{\text{layer},j} (1-u_{\text{layer},j}) u_{\text{layer-1},i};$ **endloop**

endloop

end; {COMPUTE_GRADIENT}

subroutine UPDATE_WEIGHTS

5 **for all** $w_{lj,i}$ **do**
 $w_{lj,i}(k+1) = w_{lj,i}(k) - \mu g_{lj,i};$
 endloop
end; {UPDATE_WEIGHTS}

10 A summary of the above code will now be provided.

Suppose each line of the network parameters in Fig. 2 were classified into one of ten states, where the states are labelled with numeric values, for example, 1, 2, ..., 10. (Note that this labelling is what characterizes this pseudo-code as "supervised" learning.) Then, an extra column in Fig. 2 would hold the classification of the lines of data. Each line of data in Fig. 2 may be characterized as a training pair (x,d), where x is a training sample (i.e., the original line) and d is a desired output (i.e., the classification). Given a table of training pairs, the learning algorithm proceeds as follows:

20 (i) (x,d) of the first line in the table is read;

(ii) the FEED-FORWARD subroutine uses an initial set of random weights w to map the training sample x into an actual output u;

25 (iii) the COMPUTE_GRADIENT subroutine calculates the difference between the actual output u and the desired output d; and

(iv) given this difference, the UPDATE_WEIGHTS subroutine adjusts the set of weights w so that the actual output u becomes closer to the desired output d.

30

After multiple iterations of this algorithm over the subsequent lines of the training (accumulated) data in Fig. 2, the algorithm learns to classify each line of

the training data into the correct state, and also to classify future lines of data into corresponding states. A similar algorithm may be used to learn the transitions or trends from one state to another.

This and other methods are described in the paper, "Progress in
5 Supervised Neural Networks," IEEE Signal Processing Magazine, January 1993, by Don Hush and Bill Horne. Another good summary of these methods is found in the book, Neural Network Learning, by Steve Gallant, MIT Press, 1993.

Fig. 7 shows yet another alternative embodiment of a simple
daytime/nighttime transition diagram, wherein state 1 may represent the daytime
10 state, and state 2 the nighttime state. The present method translates the average of activity on all ports of the network at any time increment into the above-defined symbolic language. Thus, if the network represented in the state transition graph of Fig. 7 is in state 1 and the average is decreasing quickly, the present method would predict that the network will enter the nighttime state.

15 Fig. 8 is a flow chart summarizing various aspects of the present invention. To the left of vertical line 50, are steps 53-55 for "learning" network behavior. To the right of line 50, are steps 56-59 for "predicting" network behavior.

In regard to learning network behavior, the method begins at step 52 and proceeds to step 53 to read a history of network parameters from the data
20 repository 18. The method then proceeds to step 54 to calculate/display a range of network states, and optionally identify the states as "bad," "good," "midday," etc. The method then proceeds to step 55 to construct/display a state transition graph as shown for example, in Fig. 5. The method may end at step 60, or proceed further as described below.

25 In regard to predicting network behavior, after completing the above steps 53-55, the method proceeds to step 56 to read current network parameters. The current network parameters are measured for a predetermined time period to determine/display the current state in step 57. Alternatively or in combination, the parameters are read for a sufficient time period to calculate/display the current trends
30 in step 58. Then, with the results of steps 55, 57 and 58, the method may then predict/display future network states in step 59.

It should be noted that the arrows in the flow chart mean that the data produced by the box at the tail end of the arrow is required to perform the task in the box at the head of the arrow, e.g., the "predict/display future network states" box 59 requires as input the state transition graph 55, the current state 57, and the current trends 58.

Generally, it is expected that the user will execute the "learning network behavior" part of the method (to the left of dotted line 50), and later run the "predicting network behavior" (on the right side of line 50). Thus, the results of steps 54 and 55 are available for steps 57 and 59 respectively. Alternatively, the "predicting network behavior" part may be run continuously.

As part of step 59 ("predict/display future network states"), a warning signal may be generated and displayed to the user to indicate that the network is moving toward a problematic state.

Having thus described various embodiments of the present invention, additional modifications and improvements will readily occur to those skilled in the art. Accordingly, the foregoing description is by way of example only, and is not intended to be limiting.

CLAIMS

1. A method of determining behavior patterns of a network from accumulated network parameters, the method including the steps of:

- 5 a) reading the accumulated network parameters;
 b) determining possible states of the network and possible transitions between the possible states from the accumulated network parameters; and
 c) displaying to a user, one or more of the possible states and the possible transitions.

10

2. The method of claim 1, further including:
 measuring current network parameters for a predetermined time period;
 determining a current state of the network from the current network parameters and the possible states; and
15 displaying to a user the current state.

15

3. The method of claim 2, further including:
 determining current trends of the network from the current network parameters and the possible transitions;
20 predicting possible next states that the network may enter from the possible states, the possible transitions, the current state, and the current trends; and
 displaying to a user at least one of the current trends and the possible next states.

20

25

4. The method of claim 3, further including:
 labeling at least one of the possible states, the possible transitions, the current state, the possible next states and the current trends according to a performance criterion.

30

- 13 -

5. The method of claim 3, further including:
designating at least one of the possible states as a problematic state;
and
warning a user when any of the possible next states is a problematic
5 state.

6. A method of determining behavior patterns of a network, the method
including the steps of:

- 10 a) measuring accumulated network parameters for a first predetermined
time period;
b) storing the accumulated network parameters in a data repository;
c) reading the accumulated network parameters from the data
repository;
d) determining possible states of the network and possible transitions
15 between the possible states from the accumulated network parameters; and
e) displaying to a user, one or more of the possible states and the
possible transitions.

7. The method of claim 6, further including:
20 measuring current network parameters for a second predetermined time
period;
determining a current state of the network from the current network
parameters and the possible states; and
displaying to a user the current state.

25

8. The method of claim 7, further including:
determining current trends of the network from the current network
parameters and the possible transitions;
predicting possible next states that the network may enter from the
30 possible states, the possible transitions, the current state, and the current trends; and
displaying to a user at least one of the current trends and the possible
next states.

9. The method of claim 8, further including:

labeling at least one of the possible states, the possible transitions, the current state, the possible next states, and the current trends according to a performance criterion.

5

10. The method of claim 9, further including:

designating at least one of the possible states as a problematic state;

and

warning a user when any of the possible next states is a problematic state.

10

11. An apparatus that reads accumulated network parameters from a data repository of a network management system, and that provides a description of behavior patterns of the network, the apparatus comprising:

15

a data reader, coupled to the data repository, that reads in the accumulated network parameters;

a behavior analyzer, coupled to the data reader, that determines possible states of the network and possible transitions between the possible states from the accumulated network parameters; and

20

a states display, coupled to the network behavior analyzer, that communicates to a user the possible states and the possible transitions.

12. The apparatus according to claim 11, further comprising:

25

a monitor, coupled to the network, that measures current network parameters for a predetermined time period;

a status analyzer, coupled to the monitor and the behavior analyzer, that determines a current state of the network from the current network parameters and the possible states; and

30

a current status display, coupled to the status analyzer, that communicates to a user the current state.

13. The apparatus according to claim 12, further comprising:

a trends analyzer, coupled to the monitor and the behavior analyzer, that determines current trends of the network from the current network parameters and the possible transitions;

5 a behavior predictor, coupled to the status analyzer, the trends analyzer, and the behavior analyzer, that determines possible next states that the network may enter from the possible states, the possible transitions, the current state, and the current trends; and

10 a predictor display, coupled to the trends analyzer and the behavior predictor, that communicates to a user at least one of the current trends and the possible next states.

14. The apparatus according to claim 13, further comprising:

15 a state identifier, coupled to at least one of the behavior analyzer, the status analyzer, and the behavior predictor, that labels at least one of the possible states, the possible transitions, the current state, the possible next states and the current trends according to a performance criterion.

15. The apparatus of claim 13, further comprising:

20 a state identifier that designates at least one of the possible states as a problematic state; and

a warning unit, coupled to the state identifier, that warns a user when any of the possible next states is a problematic state.

1/7

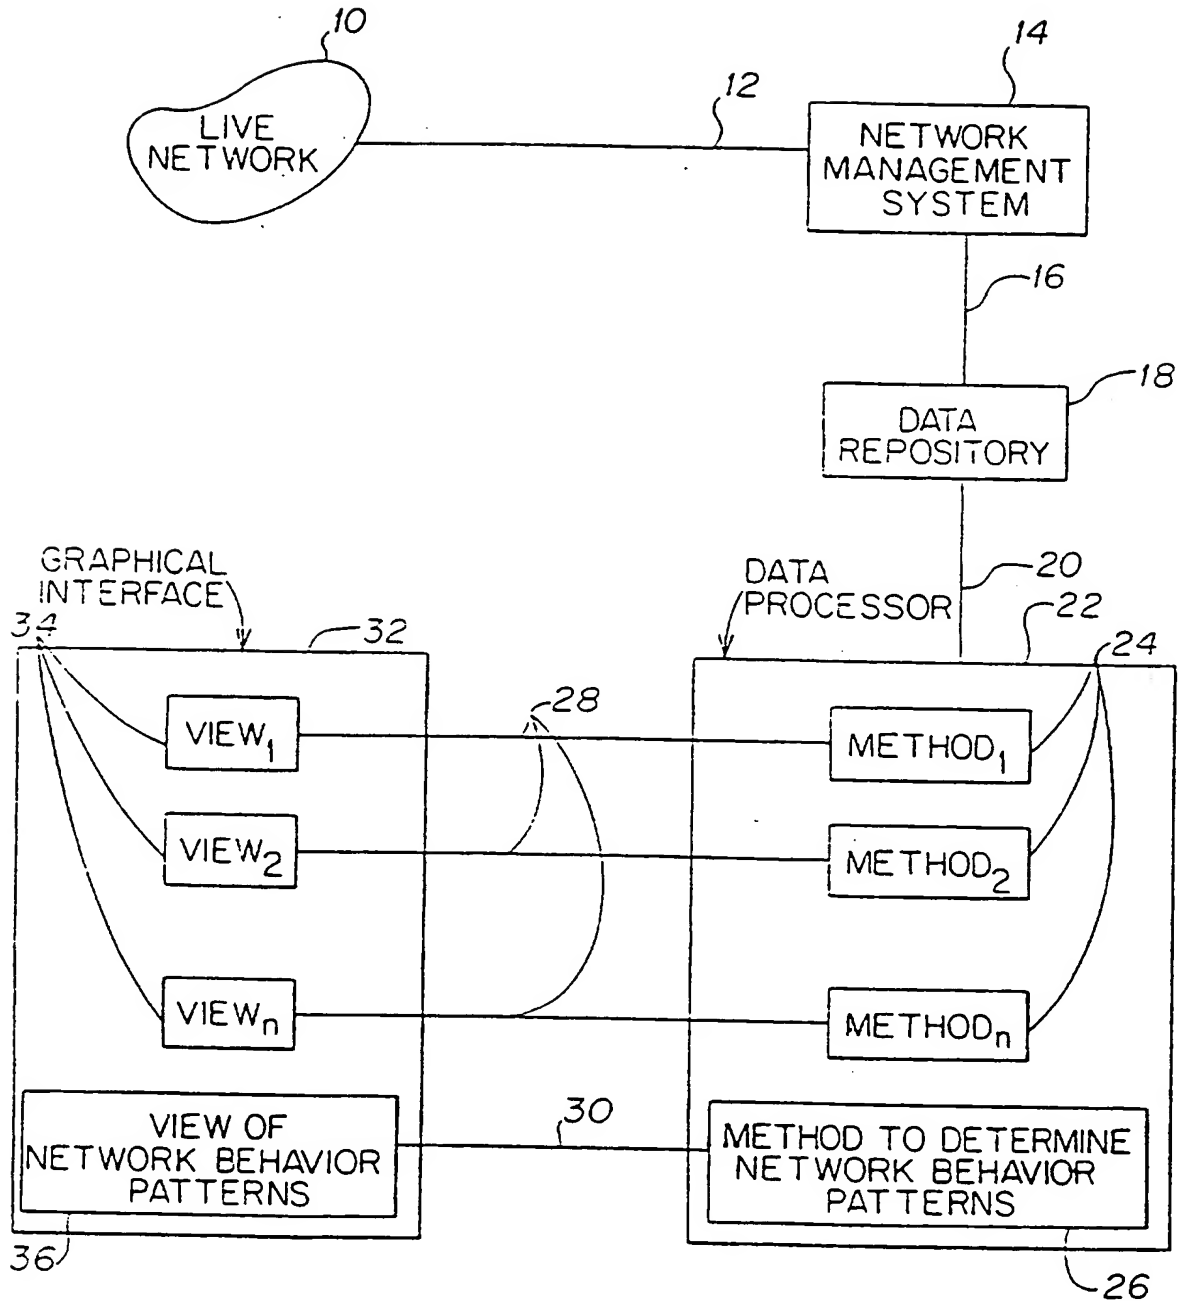


FIG. 1

2 / 7

Time	P1	P2	P3	P4	P5	P6	P7	P8	P9	P10	P11	P12
68	10.48	2.14	1.52	19.92	2.28	0.37	2.57	2.90	3.56	4.72	3.40	10.15
69	9.77	2.05	2.10	15.38	2.15	0.19	1.19	4.19	3.34	6.39	4.61	14.61
70	8.83	2.67	2.05	8.34	4.11	0.21	1.12	2.23	4.53	3.91	5.59	13.33
71	8.73	2.02	1.30	8.07	4.63	0.19	1.15	1.88	0.55	4.87	6.54	15.54
72	5.69	4.09	1.21	7.13	4.21	0.27	2.66	2.24	1.36	6.20	4.55	11.73
73	6.16	5.62	2.93	14.74	4.49	0.27	1.83	3.45	0.46	5.23	6.76	13.43
74	7.04	6.91	3.03	16.30	2.19	0.13	0.57	1.67	0.21	6.68	5.42	20.59
75	5.53	2.64	0.44	5.91	3.45	0.23	1.43	2.07	0.24	4.73	3.90	13.04
76	12.97	0.89	0.52	4.48	2.09	0.13	2.22	6.20	0.51	6.86	0.91	2.85
77	4.35	1.63	1.64	6.26	2.20	0.12	2.59	1.60	0.18	2.74	0.99	2.52
78	11.70	1.73	0.95	4.78	1.10	0.11	0.88	3.17	0.06	2.36	6.49	14.86
79	7.33	1.14	1.12	4.84	2.46	0.20	2.90	1.10	0.03	3.90	4.33	10.34

FIG. 2

3 / 7




FIG. 3
(PRIOR ART)

4/7

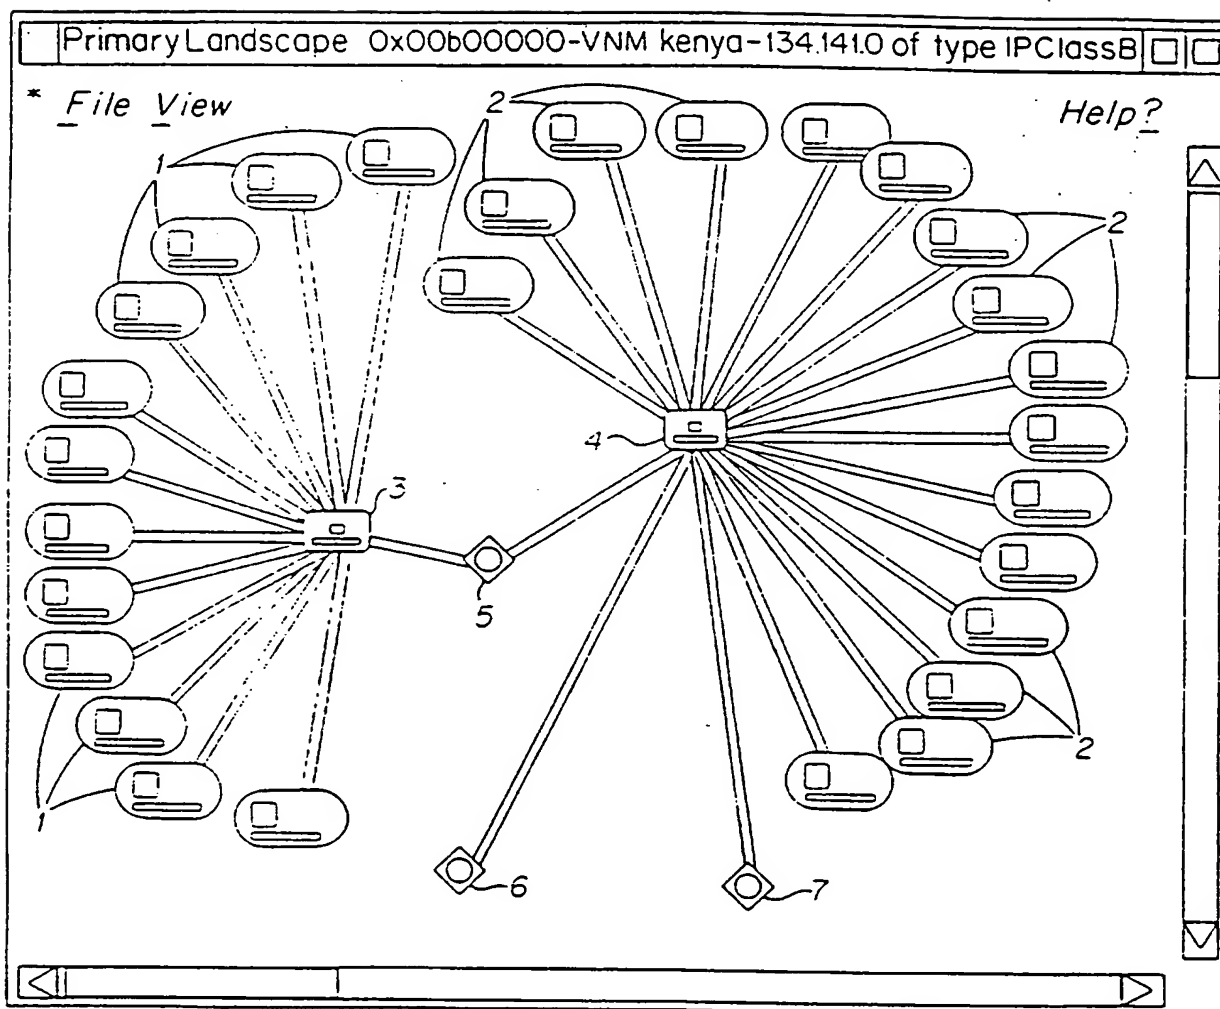


FIG. 4

5/7

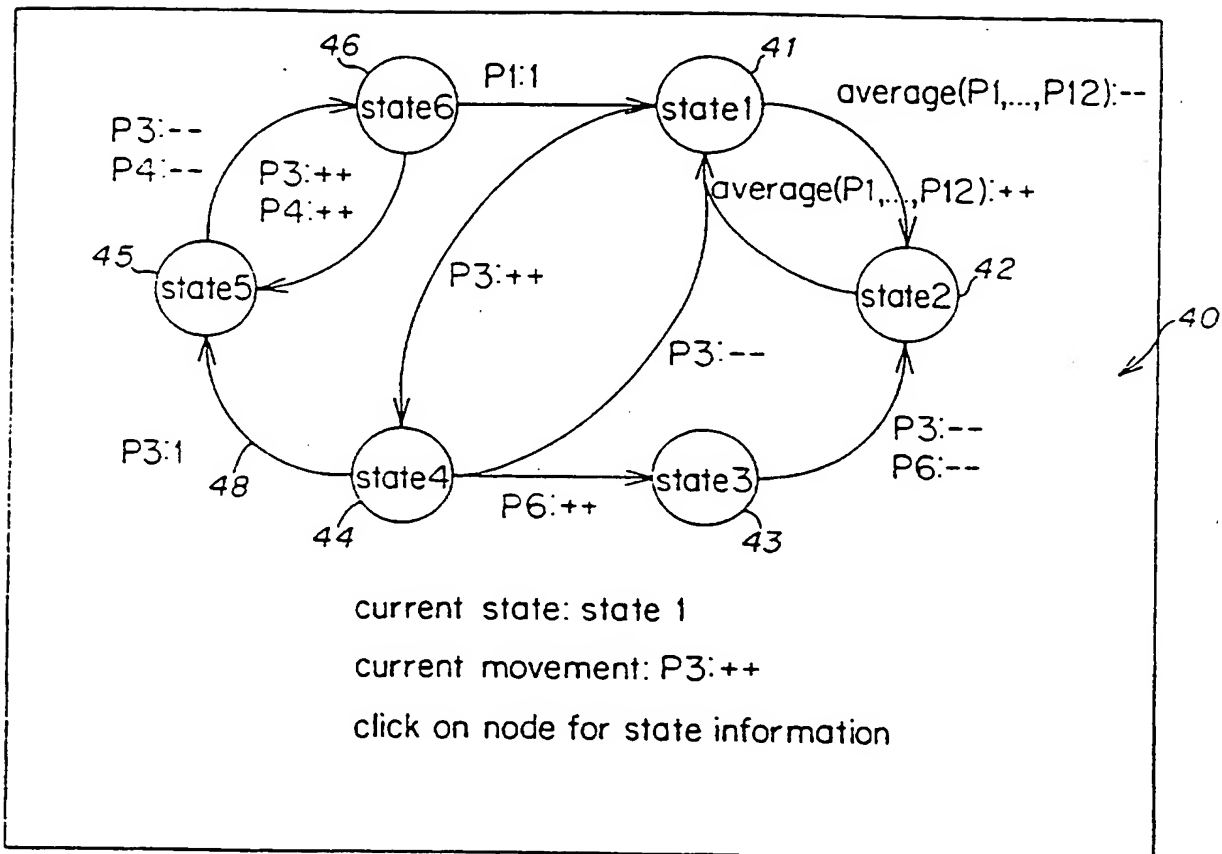


FIG. 5

6/7




FIG. 6




FIG. 7

7/7

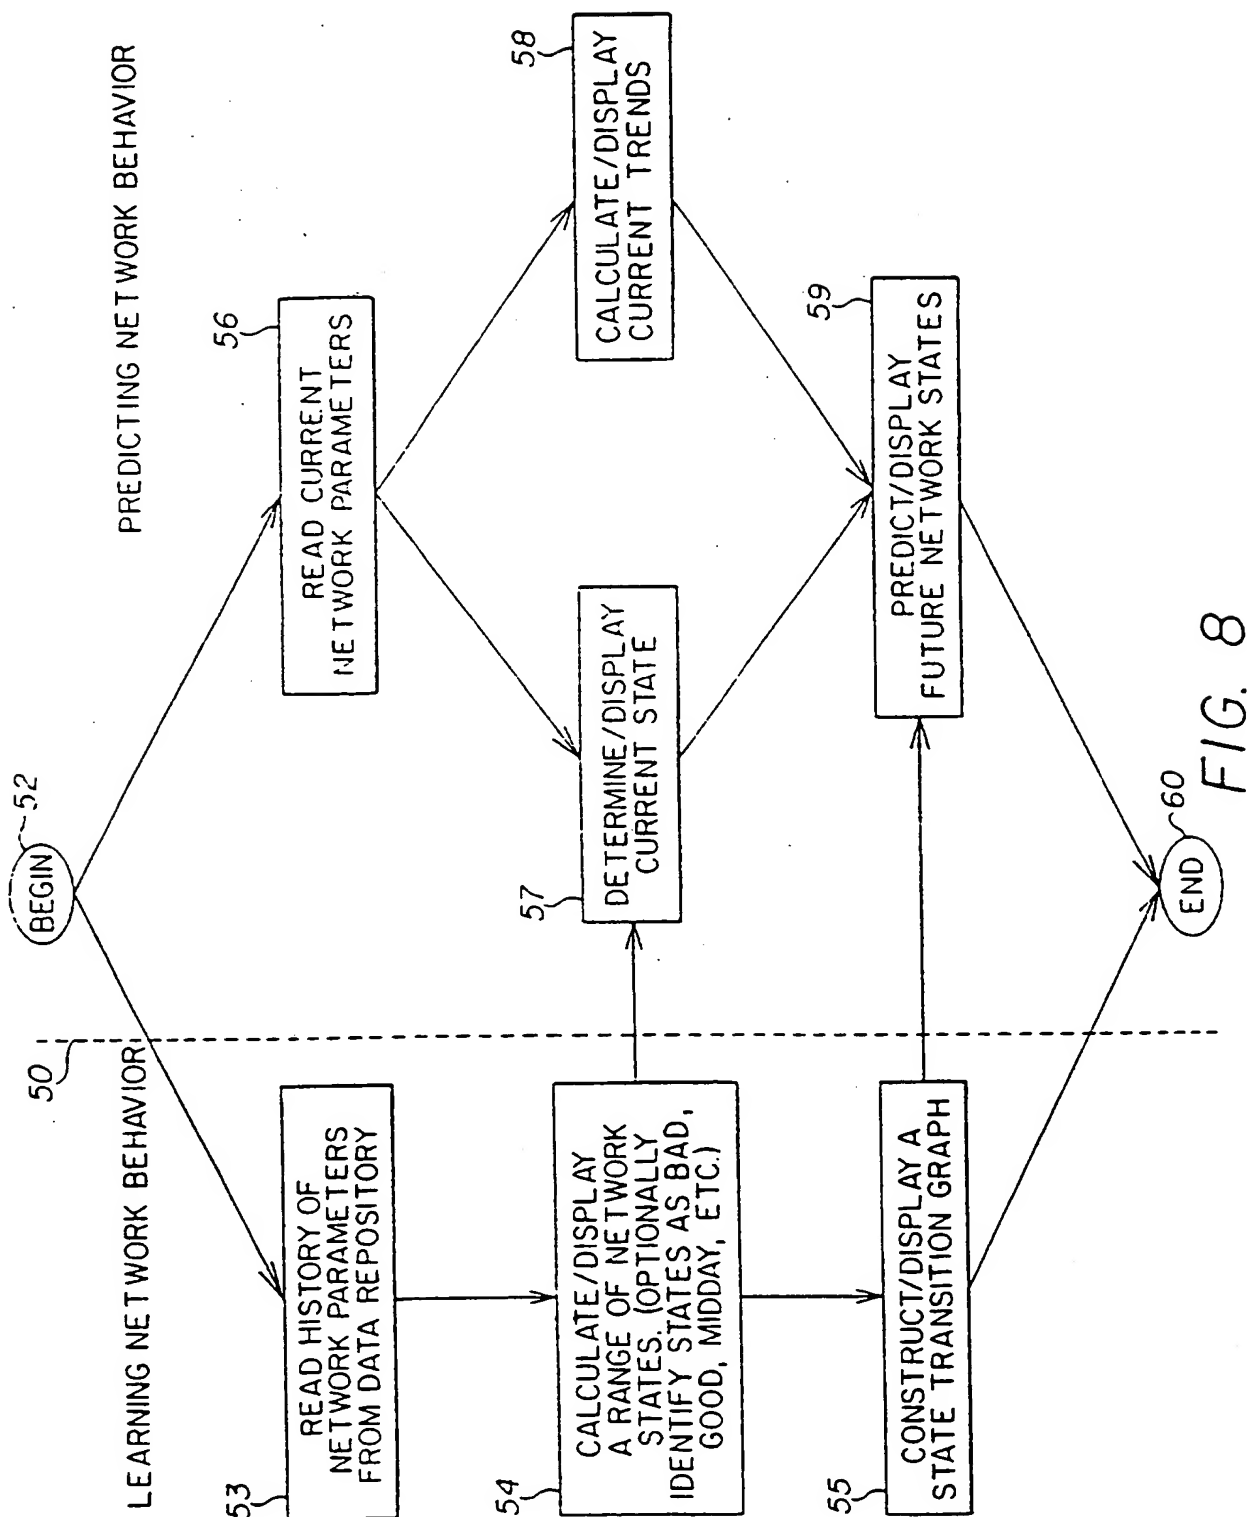


FIG. 8



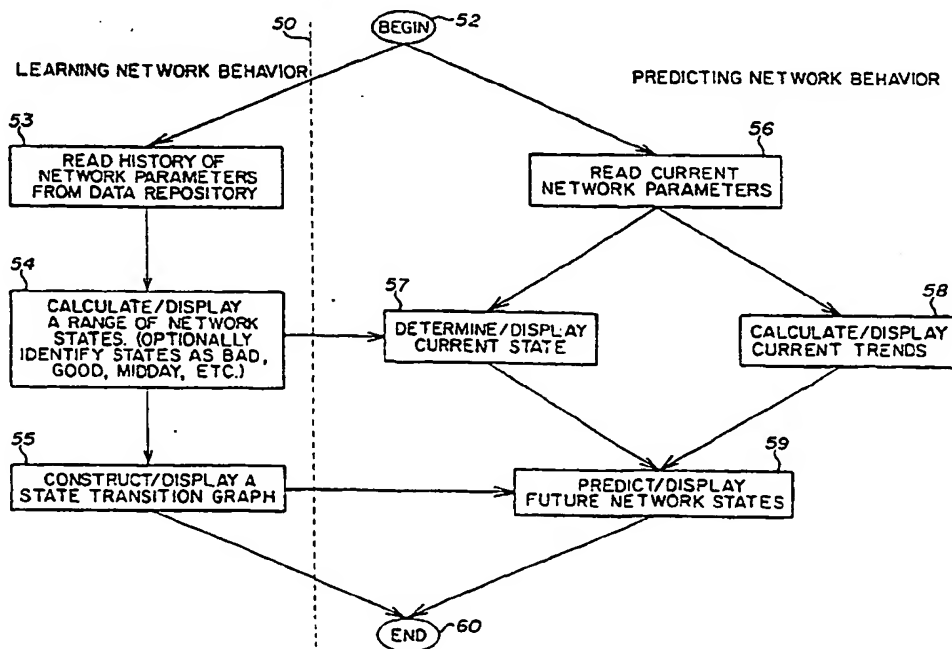
INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : H04L 12/24, 12/26		A3	(11) International Publication Number: WO 96/24210
			(43) International Publication Date: 8 August 1996 (08.08.96)
(21) International Application Number: PCT/US96/01137			(81) Designated States: AL, AM, AT, AU, AZ, BB, BG, BR, BY, CA, CH, CN, CZ, DE, DK, EE, ES, FI, GB, GE, HU, IS, JP, KE, KG, KP, KR, KZ, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, TJ, TM, TR, TT, UA, UG, UZ, VN, ARIPO patent (KE, LS, MW, SD, SZ, UG), Eurasian patent (AZ, BY, KG, KZ, RU, TJ, TM), European patent (AT, BE, CH, DE, DK, ES, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, ML, MR, NE, SN, TD, TG).
(22) International Filing Date: 24 January 1996 (24.01.96)			
(30) Priority Data: 08/382,666 2 February 1995 (02.02.95) US			
(71) Applicant: CABLETRON SYSTEMS, INC. [US/US]; 35 Industrial Way, Rochester, NH 03867 (US).			
(72) Inventors: LEWIS, Lundy; 480 Greenville Road, Mason, NH 03048 (US). DATTA, Utpal; 52 Pinecrest Drive, Bedford, NH 03110 (US).			
(74) Agent: HENDRICKS, Therese, A.; Wolf, Greenfield & Sacks, P.C., 600 Atlantic Avenue, Boston, MA 02210 (US).			<p>Published</p> <p><i>With international search report.</i></p> <p><i>Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i></p>
			(88) Date of publication of the international search report: 7 November 1996 (07.11.96)

(54) Title: METHOD AND APPARATUS FOR LEARNING NETWORK BEHAVIOR TRENDS AND PREDICTING FUTURE BEHAVIOR OF COMMUNICATIONS NETWORKS

(57) Abstract

Apparatus and method for learning current network behavior and predicting future behavior which utilizes a state transition graph (40). The graph includes nodes (41-46) which represent network states, and arcs (48) which represent trends in observable network parameters that result in a transition from a current state to another state. For example, a watch service may be instituted on multiple ports of a router (3, 4), and the observed network traffic on the ports over time may be transformed into a state transition graph that represents network behavior. The network states may be labeled such as "good", or "bad", etc., according to a predetermined performance criteria. Once a state transition graph is constructed, the system may then monitor the current state and current trends of the network parameters in order to predict and display future network states. The system may include an automatic warning signal for alerting a user that the network is headed in the direction of a problematic state.



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AM	Armenia	GB	United Kingdom	MW	Malawi
AT	Austria	GE	Georgia	MX	Mexico
AU	Australia	GN	Guinea	NE	Niger
BB	Barbados	GR	Greece	NL	Netherlands
BE	Belgium	HU	Hungary	NO	Norway
BF	Burkina Faso	IE	Ireland	NZ	New Zealand
BG	Bulgaria	IT	Italy	PL	Poland
BJ	Benin	JP	Japan	PT	Portugal
BR	Brazil	KE	Kenya	RO	Romania
BY	Belarus	KG	Kyrgyzstan	RU	Russian Federation
CA	Canada	KP	Democratic People's Republic of Korea	SD	Sudan
CF	Central African Republic	KR	Republic of Korea	SE	Sweden
CG	Congo	KZ	Kazakhstan	SG	Singapore
CH	Switzerland	LI	Liechtenstein	SI	Slovenia
CI	Côte d'Ivoire	LK	Sri Lanka	SK	Slovakia
CM	Cameroon	LR	Liberia	SN	Senegal
CN	China	LT	Lithuania	SZ	Swaziland
CS	Czechoslovakia	LU	Luxembourg	TD	Chad
CZ	Czech Republic	LV	Latvia	TG	Togo
DE	Germany	MC	Monaco	TJ	Tajikistan
DK	Denmark	MD	Republic of Moldova	TT	Trinidad and Tobago
EE	Estonia	MG	Madagascar	UA	Ukraine
ES	Spain	ML	Mali	UG	Uganda
FI	Finland	MN	Mongolia	US	United States of America
FR	France	MR	Mauritania	UZ	Uzbekistan
GA	Gabon			VN	Viet Nam

INTERNATIONAL SEARCH REPORT

Int'l Application No
PCT/US 96/01137

A. CLASSIFICATION OF SUBJECT MATTER
IPC 6 H04L12/24 H04L12/26

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 6 H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	<p>IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS, vol. 11, no. 9, December 1993, NEW YORK, NY, US, pages 1415-1425, XP000491498 YING-DAR LIN, MARIO GERLA: "Induction and Deduction for Autonomous Networks" see page 1415, left-hand column, line 1 - page 1416, left-hand column, line 12 see page 1417, left-hand column, line 13 - right-hand column, line 15 see page 1418, left-hand column, line 1 - line 17 see page 1418, right-hand column, line 33 - page 1419, left-hand column, line 20 see page 1419, right-hand column, line 2 - line 25</p> <p style="text-align: center;">--- -/--</p>	1-15

☒ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

* Special categories of cited documents:

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier document but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

- "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
- "&" document member of the same patent family

Date of the actual completion of the international search

11 September 1996

Date of mailing of the international search report

19.09.96

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+ 31-70) 340-2040, Tx. 31 651 epo nl,
Fax (+ 31-70) 340-3016

Authorized officer

Vaskimo, K

INTERNATIONAL SEARCH REPORT

Int'l Application No
PCT/US 96/01137



C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	<p>DATA COMMUNICATIONS, vol. 23, no. 16, November 1994, NEW YORK, US, pages 47-48, XP000471520 MARY JANDER: "Dynamic Duo Does Monitoring" see page 47, left-hand column, line 8 - middle column, line 2 see page 47, middle column, line 16 - right-hand column, line 9 ---</p>	1-15
A	<p>NTT REVIEW, vol. 4, no. 4, July 1992, TOKYO, JP, pages 69-74, XP000310844 JUN MATSUDA, ZENJI HAYASHI, SHUJI EDANO: "Traffic Administration Based on Daily Traffic Forecast" see page 69, left-hand column, line 20 - right-hand column, line 16 see page 70, left-hand column, line 13 - page 71, left-hand column, line 6 ---</p>	1-15
A	<p>EP,A,0 412 692 (AMERICAN TELEPHONE AND TELEGRAPH COMPANY) 13 February 1991 see column 1, line 24 - column 2, line 4 see column 2, line 51 - column 3, line 35 see column 16, line 38 - column 17, line 14 ---</p>	1-4,6-9, 11-14
A	<p>DATA COMMUNICATIONS, vol. 21, no. 6, April 1992, NEW YORK, US, pages 43-44, XP000271117 JIM CARR: "Low-Cost Network Monitor Gets the LAN Lowdown" see page 43, left-hand column, line 32 - line 39 see page 43, right-hand column, line 5 - line 19 -----</p>	1-3,6-8, 11-13

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/US 96/01137

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
EP-A-412692	13-02-91	CA-A- 2017974	07-02-91
		DE-D- 69013269	17-11-94
		DE-T- 69013269	24-05-95
		ES-T- 2060960	01-12-94
		JP-A- 3071193	26-03-91
		JP-B- 7036106	19-04-95
		US-A- 5136690	04-08-92
